arXiv:1703.06206v1 [stat.CO] 17 Mar 2017

Sequential Monte Carlo Methods in the **nimble** R Package

Nicholas Michaud[1,2], Perry de Valpine[2], Daniel Turek[3], Christopher J. Paciorek[1]

[1]Department of Statistics, University of California, Berkeley

[2]Department of Environmental Science, Policy, and Management, University of California, Berkeley

[3] Department of Mathematics and Statistics, Williams College

Correspondence to: Nicholas Michaud, nmichaud@berkeley.edu

Abstract

nimble is an R package for constructing algorithms and conducting inference on hierarchical models. The **nimble** package provides a unique combination of flexible model specification and the ability to program model-generic algorithms – specifically, the package allows users to code models in the BUGS language, and it allows users to write algorithms that can be applied to any appropriately-specified BUGS model. In this paper, we introduce **nimble**'s capabilities for state-space model analysis using Sequential Monte Carlo (SMC) techniques. We first provide an overview of state-space models and commonly used SMC algorithms. We then describe how to build a state-space model and conduct inference using existing SMC algorithms within **nimble**. SMC algorithms within **nimble** currently include the bootstrap filter, auxiliary particle filter, Liu and West filter, ensemble Kalman filter, and a particle MCMC sampler. These algorithms can be run in R or compiled into C++ for more efficient execution. Examples of applying SMC algorithms to a random walk model and a stochastic volatility model are provided. Finally, we give an overview of how model-generic algorithms are coded within **nimble** by providing code for a simple SMC algorithm.

1 Introduction

State-space models provide a method for analyzing time series data, where observations are assumed to be noisy measurements of unobserved latent states that evolve over time. State-space models have been used in such diverse fields as population ecology (Knape and de Valpine 2012), epidemiology (Andersson and Britton 2000), economics (Creal 2012), and meteorology (Wikle *et al.* 2013).

With the broad applicability of state-space models has come a variety of techniques for conducting inference. One common goal of inference for state-space models is determining the filtering distribution of the model, that is, the distribution of the latent states given data up to a certain time point. A second goal lies in estimating the likelihood of the model by integrating over the latent state dimensions. For state-space models that do not follow a linear Gaussian framework, analytical solutions to the filtering distribution and likelihood problems are usually unavailable. For such models, inference and estimation is commonly performed using a set of flexible computational algorithms known as Sequential Monte Carlo (SMC) methods (Doucet *et al.* 2001).

SMC methods are attractive as they provide a general framework for conducting inference on any state-space model. In addition, SMC methods generally perform "on-line" inference, that is, inference on filtering distributions that can be updated sequentially as more data are received. A variety of SMC methods currently exist, including the bootstrap filter (Gordon *et al.* 1993), auxiliary particle filter (Pitt and Shephard 1999a), Liu and West filter (Liu and West 2001), Storvik filter (Storvik 2002), particle learning algorithm (Carvalho *et al.* 2010), and others. In addition, algorithms such as Particle MCMC (Andrieu *et al.* 2010) have been developed that place SMC methods within a broader MCMC framework. Although the different SMC algorithms estimate the filtering distribution using a variety of techniques, they are tied together through their use of sequential importance resampling (Doucet *et al.* 2001) to update filtering estimates as new data are received.

The generality with which SMC methods can be applied makes them well suited for use within the **nimble** R software package. **nimble** (de Valpine *et al.* in-press) allows hierarchical models to be written using the BUGS language. These models can then be analyzed using **nimble**'s library of model-generic algorithms. Additionally, **nimble** provides a domain specific language (DSL) with which users can write their own model-generic algorithms. These algorithms can be run in R or compiled into C++ for more efficient execution. In this paper, **nimble**'s SMC algorithms are described in detail. **nimble** also has a variety of MCMC algorithms for more general Bayesian inference, as well as an MCEM algorithm.

Other current software packages that implement SMC algorithms include the **pomp** R package (King *et al.* 2016), the **LibBi** package (Murray 2015), the **Biips** package (Todeschini *et al.* 2014), and the **vSMTC** C++ template library (Zhou 2015). **nimble** differs from the aforementioned software in its focus on providing an accessible DSL for writing algorithms that can be applied to any model written in the BUGS language. As such, the **nimble** software package offers a useful platform for extending existing SMC algorithms or programming new ones.

In Section 2, we introduce state-space models and the idea of filtering distributions. We then describe a variety of algorithms that can be used for inference on filtering distributions. In Section 3, we provide examples of specifying state-space models within **nimble**. Inference is conducted using **nimble**'s SMC algorithms and the results are analyzed. In Section 4, we demonstrate **nimble**'s programmability by coding an SMC algorithm within nimble's DSL.

2 Sequential Monte Carlo methods for state-space models

2.1 State-space models

State-space models, also known as hidden Markov models, are used to model time series data or any data that arrives sequentially. The vector of data at each time t, labeled y_t, is assumed to be related to a latent, unobserved state x_t through an observation equation $y_t \sim g_t(y_t|x_t, \theta)$. Here, θ is a vector of top-level parameters that are assumed not to change with time. In addition to the observation equation describing the dependence of y_t on x_t, x_t depends on x_{t-1} through a transition equation $x_t \sim f_t(x_t|x_{t-1}, \theta)$. Both the observation and transition equation are stochastic. Frequently, the observation and transition equations remain constant over all time points, in which case the t subscripts on f_t and g_t are dropped. State-space models have the following conditional independence property: $[x_t|x_{1:t-1}] = [x_t|x_{t-1}]$, where $x_{1:t-1} = (x_1, \ldots, x_{t-1})$. Note that we assume no observation exists for $t = 0$, and that x_0 comes from a known prior distribution $p_0(x_0|\theta)$.

One common inferential goal when analyzing state-space models lies in determining the distribution $p(x_t|y_{1:t}, \theta)$, known as the filtering distribution for x_t. Consider a situation where new data are received sequentially in time, and data are currently available up to time $t-1$ – that is, $y_{1:t-1}$ is known. Upon receiving y_t, the filtering distribution $p(x_t|y_{1:t}, \theta)$ provides information about the latent state at the most recent time point, given the most recent data. Another common goal of inference is calculating the likelihood $p(y_{1:t}|\theta)$. This likelihood can in turn be used for model selection or in an MCMC framework to obtain samples from $p(\theta|y_{1:t})$.

For certain types of state-space models, the filtering distributions and likelihood are analytically tractable. For example, the Kalman filter (Kalman 1960) can be used to derive the filtering distribution for state-space models in which both the observation and transition equations are linear and have a Gaussian error term. However, there are a wide variety of state-space models for which no analytical solutions are available.

Two estimation methods are common for state-space models with intractable filtering distributions and likelihoods. MCMC algorithms can be used to draw samples from $p(x_{1:t}|y_{1:t}, \theta)$, if the top-level parameters are assumed to be fixed, or from $p(x_{1:t}, \theta|y_{1:t})$ if they are not fixed. A second group of methods for estimating the filtering distribution for non-linear or non-Gaussian state-space models are known as Sequential Monte Carlo (SMC) methods, or particle filters.

2.2 Filtering algorithms

In Section 2.3 and Section 2.4, two types of SMC methods (the bootstrap filter and auxiliary particle filter) are described, each of which can be used to generate samples from the filtering distribution $p(x_t|y_{1:t}, \theta)$ or the smoothing distribution $p(x_{1:t}|y_{1:t}, \theta)$. A third method, known as the Liu and West filter, can be used

to concurrently sample values of any top-level parameters along with latent states, resulting in an approximation of $p(x_t, \theta|y_{1:t})$. The Liu and West filter is described in Section 2.5. In Section 2.6, a particle MCMC algorithm is detailed that uses particle filters to estimate likelihoods within a Metropolis-Hastings MCMC sampling scheme for θ.

The Ensemble Kalman filter, or EnKF (Evensen 2003), can also be used to conduct inference on the filtering distribution of latent states. Similar to SMC techniques, the EnKF approximates the filtering distribution via a collection of particles that are propagated forwards in time. However, whereas SMC methods use importance sampling to select particles at each time point, the EnKF instead shifts particles towards the filtering distribution using an approximation to the Kalman gain matrix. The EnKF is described in Section 2.7.

2.3 Bootstrap filter

The bootstrap filter of Gordon *et al.* (1993) uses importance sampling to sequentially generate samples from $p(x_t|y_{1:t})$ at each time t. Note that since the bootstrap filter operates for given values of top level parameters, we omit the dependence on θ in our notation. Specifically, suppose that we have K samples, called particles, from $p(x_{t-1}|y_{t-1})$ labeled $x_{t-1}^{(k)}$ for $k = 1, \ldots, K$. The bootstrap filter first propagates each of these particles forward according to a proposal distribution $\tilde{x}_t^{(k)} \sim q(x_t|x_{t-1}^{(k)}, y_t)$. Importance weights $\pi_t^{(k)}$ are then calculated for each particle, and the propagated particles are resampled according to these weights. This results in an equally weighted sample $(x_t^{(k)})_{k=1}^K$ from $p(x_t|y_{1:t})$.

Algorithm 1 Bootstrap filter

1: **for** k in $1:K$ **do**
2: Generate $x_0^{(k)} \sim p_0(x_0)$
3: Set $\pi_0^{(k)} = \frac{1}{K}$
4: **end for**
5: **for** t in $1:T$ **do**
6: **for** k in $1:K$ **do**
7: Generate $\tilde{x}_t^{(k)} \sim q(x_t|x_{t-1}^{(k)}, y_t)$
8: Calculate $w_t^{(k)} = \frac{f(\tilde{x}_t^{(k)}|x_{t-1}^{(k)}) g(y_t|\tilde{x}_t^{(k)})}{q(\tilde{x}_t^{(k)}|x_{t-1}^{(k)}, y_t)} \pi_t^{(k)}$
9: Normalize $w_t^{(k)}$ as $\pi_t^{(k)} = \frac{w_t^{(k)}}{\sum_{i=1}^K w_t^{(i)}}$
10: Sample $x_t^{(k)} \sim \sum_{i=1}^K \pi_t^{(i)} \delta(x - \tilde{x}_t^{(i)})$
11: Set $\pi_t^{(k)} = \frac{1}{K}$
12: **end for**
13: Calculate $\tilde{p}(y_{t|1:t-1}) = \frac{1}{K} \sum_{k=1}^K w_t^{(k)}$
14: **end for**

Note that resampling (Step 10 in Algorithm 1) creates an equally weighted sample from the target distribution. In the resampling step, δ is defined as the

Dirac delta function. Additionally, an estimate of the likelihood $p(y_{1:T})$ can be obtained by $\tilde{p}(y_{1:T}) = \prod_{t=1}^{T} \tilde{p}(y_{t|1:t-1})$, where $\tilde{p}(y_{t|1:t-1})$ is given in line 13 of the algorithm.

The resampling step of Algorithm 1 is performed to reduce particle degeneracy. Particle degeneracy is a commonly occurring problem in filtering algorithms, where a small number of particles have most of the weight placed on them, while the majority of particles have practically zero weight (Doucet *et al.* 2000). Particle degeneracy corresponds to high Monte Carlo variance of approximations made using the filtered particles, and causes the filter to spend computational effort in propagating and weighting particles that contribute little to our knowledge of the target distribution. Resampling ensures that mostly highly-weighted particles will be propagated forwards, increasing algorithm efficiency and providing a better estimate of the target distribution.

However, resampling particles at each time point can lead to a loss of particle "diversity" (Doucet *et al.* 2000), as many of the resampled particles at each time point will have the same value. Thus it has been proposed (Smith *et al.* 2001) that resampling should take place only if particle degeneracy becomes too significant. An estimate of particle degeneracy is the effective sample size, calculated at each time t as $ESS = \frac{1}{\sum_{k=1}^{K} \left(\pi_t^{(k)}\right)^2}$. Smith *et al.* (2001) recommend that resampling should be conducted only if the effective sample size becomes too low, indicating many particles with low weights. As a criterion for when a resampling step should take place, a threshold τ must be chosen with $0 \leq \tau \leq 1$, such that the algorithm will resample particles whenever $\frac{ESS}{K} < \tau$. Note that choosing $\tau = 0$ will mean that the resampling step is never performed, and choosing $\tau = 1$ will ensure that sampling is performed at each time point. To perform the above algorithm without resampling, simply remove Steps 10 and 11. If the resampling step is not performed, the set $(\tilde{x}_t^{(k)}, \pi_t^{(k)})$ will constitute a non-equally weighted sample from the target distribution. Various methods for resampling particles have been employed, including systematic resampling, residual resampling, and multinomial resampling (Doucet and Johansen 2009). Additionally, the above filter can be used to produce samples from the smoothing distribution $p(x_{1:t}|y_{1:t})$ using methods described in Doucet and Johansen (2009).

2.4 Auxiliary particle filter

The auxiliary particle filter algorithm (APF) of Pitt and Shephard (1999a) uses importance sampling similarly to the bootstrap filter but includes an additional "look-ahead step". At each time point t, the auxiliary particle filter algorithm calculates first-stage weights $w_{t|t-1}^{(k)}$ for particles from time $t-1$. These weights are calculated using an estimate of the likelihood of the current data given each particle from the previous time point, labeled $\hat{p}(y_t|x_{t-1}^{(k)})$. Particles with high first-stage weights correspond to values of the latent state at time $t-1$ that are likely to generate the observed data at time t. To estimate $\hat{p}(y_t|x_{t-1}^{(k)})$ Pitt and Shephard (1999a) recommend choosing an auxiliary variable $\tilde{x}_{t|t-1}^{(k)}$ and then

setting $\hat{p}(y_t|x_{t-1}^{(k)}) = g(y_t|\tilde{x}_{t|t-1}^{(k)})$. Possible methods for choosing $\tilde{x}_{t|t-1}^{(k)}$ include simulating a value from $f(x_t|x_{t-1}^{(k)})$, or taking $\tilde{x}_{t|t-1}^{(k)} = E(x_t|x_{t-1}^{(k)})$.

The first-stage weights are used to sample K particles from time $t-1$, labeled $\tilde{x}_{t-1}^{(k)}$ for $k = 1, \ldots, K$. The sampled particles are then propagated forwards by a proposal distribution $q(x_t^{(k)}|\tilde{x}_{t-1}^{(k)}, y_t)$ and reweighted using second-stage weights $w_t^{(k)}$, providing a weighted sample from $p(x_t|y_{1:t})$. The APF as shown in Pitt and Shephard (1999a) optionally includes a second resampling step after Step 12, using the second-stage weights. However, the algorithm using a single resampling step has been shown to be more efficient (Carpenter *et al.* 1999).

Algorithm 2 Auxiliary particle filter

1: **for** k in $1:K$ **do**
2: Generate $x_0^{(k)} \sim p_0(x_0)$
3: Set $\pi_0^{(k)} = \frac{1}{K}$
4: **end for**
5: **for** t in $1:T$ **do**
6: **for** k in $1:K$ **do**
7: Compute $w_{t|t-1}^{(k)} = \pi_{t-1}^{(k)} \hat{p}(y_t|x_{t-1}^{(k)})$
8: Normalize $w_{t|t-1}^{(k)}$ as $\pi_{t|t-1}^{(k)} = \frac{w_{t|t-1}^{(k)}}{\sum_{i=1}^{K} w_{t|t-1}^{(i)}}$
9: Sample $\tilde{x}_{t-1}^{(k)} \sim \sum_{i=1}^{K} \pi_{t|t-1}^{(i)} \delta(x - x_{t-1}^{(i)})$
10: Sample $x_t^{(k)} \sim q(x_t|\tilde{x}_{t-1}^{(k)}, y_t)$
11: Calculate $w_t^{(k)} = \frac{f(x_t^{(k)}|\tilde{x}_{t-1}^{(k)}) g(y_t|x_t^{(k)})}{\hat{p}(y_t|\tilde{x}_{t-1}^{(k)}) q(x_t^{(k)}|\tilde{x}_{t-1}^{(k)}, y_t)}$
12: Normalize $w_t^{(k)}$ as $\pi_t^{(k)} = \frac{w_t^{(k)}}{\sum_{i=1}^{K} w_t^{(i)}}$
13: **end for**
14: Calculate $\tilde{p}(y_{t|1:t-1}) = \left(\sum_{k=1}^{K} \frac{w_t^{(k)}}{K} \right) \left(\sum_{k=1}^{K} w_{t|t-1}^{(k)} \right)$
15: **end for**

In a manner similar to the bootstrap filter, the APF can be used to obtain an estimate of the likelihood $p(y_{1:T})$ as $\tilde{p}(y_{1:T}) = \prod_{t=1}^{T} \tilde{p}(y_{t|1:t-1})$, where $\tilde{p}(y_{t|1:t-1})$ is given in line 14 of the APF algorithm.

2.5 Liu and West filter

Unlike the bootstrap and auxiliary particle filters, the Liu and West filter (Liu and West 2001) allows inference to be conducted on both the latent states and the fixed parameters. Although variations on the Liu and West filter with the potential for increased efficiency have been proposed (Polson *et al.* 2008), we present the original filter. At each time point t, the Liu and West filter provides samples from $p(x_t, \theta|y_{1:t})$, the joint posterior distribution of the latent states and fixed parameters.

Suppose we have a sample of K particles from $p(x_{t-1}, \theta|y_{1:t-1})$, labeled $(x_{t-1}^{(k)}, \theta_{t-1}^{(k)})$ for $k = 1, \ldots, K$. Note that θ_{t-1} is not meant to imply that the θ parameters vary over time, as they are fixed parameters, but is rather a notation to denote our estimates of θ at time $t-1$. The Liu and West filter proceeds by first calculating first-stage weights for both the latent states and fixed parameters in a manner similar to the auxiliary particle filter. For each particle, an auxiliary value for the latent state is calculated and labeled $x_{t|t-1}^{(k)}$. Liu and West choose $x_{t|t-1}^{(k)} = E(x_t|x_{t-1}^{(k)}, \theta_{t-1}^{(k)})$. An auxiliary value for the fixed parameters is also calculated as $\theta_{t|t-1}^{(k)} = a\theta_{t-1}^{(k)} + (1-a)\bar{\theta}_{t-1}$, where $\bar{\theta}_{t-1}$ is the average of all θ particles at stage $t-1$, and where a is a known shrinkage coefficient. These auxiliary values $(x_{t|t-1}^{(k)}, \theta_{t|t-1}^{(k)})$ are then given first-stage weights $w_{t|t-1}^{(k)}$, and indices j_k are sampled using these weights for $k = 1, \ldots, K$.

After the first-stage indices are sampled, fixed parameters $\theta_t^{(k)}$ are drawn from a normal density with mean $\theta_{t|t-1}^{(j_k)}$ and variance $h^2 V_{t-1}$, where h^2 is a scaling parameter and V_{t-1} is the covariance matrix of the parameter particles $\theta_{t-1}^{(j_k)}$ with weights $\pi_{t-1}^{(j_k)}$ from the previous time point. Finally, state particles are propagated forwards and given second stage weights $w_t^{(k)}$.

Both the shrinkage coefficient a and the scaling parameter h^2 are governed by a discount parameter $d \in (0, 1]$, where $a = \frac{3d-1}{2d}$ and $h^2 = 1 - a^2$. Liu and West (2001) recommend setting d within 0.95 and 0.99.

Algorithm 3 Liu and West filter

for k in $1:K$ **do**
  Generate $x_0^{(k)} \sim p_0(x_0)$
  Generate $\theta_0^{(k)} \sim p_0(\theta)$
  Set $\pi_0^{(k)} = \frac{1}{K}$
end for
for t in $1:T$ **do**
  for k in $1:K$ **do**
    Compute $x_{t|t-1}^{(k)} = E(x_t|x_{t-1}^{(k)}, \theta_{t-1}^{(k)})$
    Compute $\theta_{t|t-1}^{(k)} = a\theta_{t-1}^{(k)} + (1-a)\bar{\theta}_{t-1}$
    Compute $w_{t|t-1}^{(k)} = \pi_{t-1}^{(k)} g(y_t|x_{t|t-1}^{(k)}, \theta_{t|t-1}^{(k)})$
    Normalize $w_{t|t-1}^{(k)}$ as $\pi_{t|t-1}^{(k)} = \frac{w_{t|t-1}^{(k)}}{\sum_{i=1}^{K} w_{t|t-1}^{(i)}}$
    Sample an index j_k from the set $\{1, \ldots, K\}$ with probabilities $\{\pi_{t|t-1}^{(i)}\}_{i=1}^{K}$
    Sample $\theta_t^{(k)} \sim N(\theta_{t|t-1}^{(j_k)}, h^2 V_{t-1})$
    Sample $x_t^{(k)} \sim f(x_t|x_{t-1}^{(j_k)}, \theta_t^{(k)})$
    Calculate $w_t^{(k)} = \frac{g(y_t|x_t^{(k)}, \theta_t^{(k)})}{g(y_t|x_{t|t-1}^{(j_k)}, \theta_{t|t-1}^{(j_k)})}$
    Normalize $w_t^{(k)}$ as $\pi_t^{(k)} = \frac{w_t^{(k)}}{\sum_{i=1}^{K} w_t^{(i)}}$
  end for
end for

2.6 Particle MCMC methods

Particle MCMC methods (Andrieu *et al.* 2010) also allow joint sampling from the posterior distribution of the states and the fixed parameters but in a manner quite different from the Liu and West filter. Particle MCMC takes advantage of the ability of certain particle filters to provide estimates of the marginal likelihood of the data, that is, $\tilde{p}(y_{1:T}|\theta) \approx \int_X p(y_{1:T}|x_{1:T}, \theta) p(x_{1:T}|\theta) dx_{1:T}$. For example, the bootstrap filter and auxiliary filter (Pitt 2002) can both be used to provide unbiased estimates of the marginal likelihood, as detailed in Sections 2.3 and 2.4. Below, we detail the Particle Marginal Metropolis Hastings (PMMH) algorithm, one of three algorithms provided in Andrieu *et al.* (2010). The PMMH algorithm is available as a sampler within **nimble**'s MCMC framework, providing a generally applicable and highly customizable method for inference.

At each iteration i, the PMMH algorithm first proposes a value θ^* for the model parameters θ from a proposal distribution $q(\theta^*|\theta^{i-1})$. Using this proposed value for θ, a particle filter is then run, which provides an estimate $\tilde{p}(y_{1:T}|\theta^*)$ of the likelihood given the proposed parameters. This marginal likelihood estimate is used to generate a sample via a Metropolis-Hastings sampler. Each iteration

of the algorithm will provide us with a single sample from $p(\theta|y_{1:T})$. If interest also lies in getting samples from the filtering distribution of the latent states, an index l can be sampled from $\{1, \ldots, K\}$ using the particle filter weights $\{\pi_t^{(k)}\}_{k=1}^K$ at each iteration i, and the particle chain $x_{1:T}^{(l)}$ can be drawn from the output of the particle filter at that iteration.

Algorithm 4 PMMH algorithm

1: Choose an initial value θ^0
2: Run an SMC algorithm to get a sample $x_{1:T}^0 \sim p(x_{1:T}|y_{1:T}, \theta^0)$ and a marginal likelihood estimate $\tilde{p}(y_{1:T}|\theta^0)$
3: **for** iteration $i \geq 1$ **do**
4: Sample $\theta^* \sim q(\theta|\theta^{i-1})$
5: Run an SMC algorithm to get a sample $x_{1:T}^* \sim p(x_{1:T}|y_{1:T}, \theta^*)$ and a marginal likelihood estimate $\tilde{p}(y_{1:T}|\theta^*)$
6: Compute $a^* = 1 \wedge \frac{\tilde{p}(y_{1:T}|\theta^*)p(\theta^*)}{\tilde{p}(y_{1:T}|\theta^{t-1})p(\theta^{t-1})} \frac{q(\theta^{t-1}|\theta^*)}{q(\theta^*|\theta^{t-1})}$
7: Generate $r \sim unif(0, 1)$
8: **if** $a^* > r$ **then**
9: Set $\theta^i = \theta^*$ and $x_{1:T}^i = x_{1:T}^*$
10: **else**
11: Set $\theta^i = \theta^{i-1}$ and $x_{1:T}^i = x_{1:T}^{i-1}$
12: **end if**
13: **end for**

2.7 Ensemble Kalman filter

In Section 2.2, the Kalman filter was mentioned as providing an analytic solution to the filtering problem when working with a linear, Gaussian state-space model. When using a model with non-linear transition equations or observation equations, however, the Kalman filter is no longer applicable. One approximation to the filtering problem for Gaussian state-space models with non-linear transition or observation equations is the ensemble Kalman filter (EnKF), which uses a particle representation of the latent states at each time point. Although the EnKF's particle representation mirrors that of the bootstrap and auxiliary particle filters described in Sections 2.3 and 2.4, the EnKF updates the latent state particles using a fundamentally different approach than the SMC methods described previously. Instead of using a sequential importance resampling framework, the EnKF first propagates particles forward using the transition equation, and then approximates the filtering distribution at each time using a multivariate normal distribution. An overview of the EnKF can be found in Gillijns *et al.* (2006). In addition, Evensen (2003) provides a comprehensive list of papers that either use or propose modifications to the EnKF.

The EnKF assumes the following forms for the observation and transition equations:

$$x_t = f(x_{t-1}) + w_t \tag{1}$$
$$y_t = g(x_t) + v_t \tag{2}$$

where w_t and v_t are normally distributed error terms with covariance matrices Q_t and R_t respectively. At each time t, assume that we have a sample of K particles from $p(x_{t-1}|y_{t-1})$ labeled $x_{t-1}^{(k)}$ for $k = 1, \ldots, K$. The particles are propagated forward according to Equation 1, giving a sample $\tilde{x}_t^{(k)}$. From these particles, a length K vector of latent errors $e_t^x = (\tilde{x}_t^{(1)} - \bar{\tilde{x}}_t, \ldots, \tilde{x}_t^{(K)} - \bar{\tilde{x}}_t)$ is calculated, where $\bar{\tilde{x}}_t$ is the average latent state taken over all particles at time t. Additionally, a vector of observation errors is calculated as $e_t^y = (\tilde{y}_t^{(1)} - \bar{\tilde{y}}_t, \ldots, \tilde{y}_t^{(K)} - \bar{\tilde{y}}_t))$, where $\tilde{y}_t^{(k)} = g(\tilde{x}_t^{(k)})$. From these error vectors, an approximate Kalman gain matrix $\tilde{K}_t$ is calculated, which in turn is used to adjust the $\tilde{x}_t^{(k)}$ particles to provide a sample from $p(x_t|y_{1:t})$.

Algorithm 5 Ensemble Kalman filter

for k in $1:K$ **do**
  Generate $x_0^{(k)} \sim p_0(x_0)$
end for
for t in $1:T$ **do**
  for k in $1:K$ **do**
    Generate $\tilde{x}_t^{(k)} \sim p(\tilde{x}_t|x_{t-1}^{(k)})$
    Calculate $\tilde{y}_t^{(k)} = g(\tilde{x}_t^{(k)})$
  end for
  Calculate $e_t^x = (\tilde{x}_t^{(1)} - \bar{\tilde{x}}_t, \ldots, \tilde{x}_t^{(K)} - \bar{\tilde{x}}_t)$
  Calculate $e_t^y = (\tilde{y}_t^{(1)} - \bar{\tilde{y}}_t, \ldots, \tilde{y}_t^{(K)} - \bar{\tilde{y}}_t)$
  Calculate $\tilde{P}_t^{xy} = \frac{1}{K-1} e_t^x (e_t^y)'$
  Calculate $\tilde{P}_t^{yy} = \frac{1}{K-1} e_t^y (e_t^y)'$
  Calculate $\tilde{K}_t = \tilde{P}_t^{xy} (\tilde{P}_t^{yy})^{-1}$
  for k in $1:K$ **do**
    Generate $v_t^{(k)} \sim N(0, R_t)$
    Calculate $x_t^{(k)} = \tilde{x}_t^{(k)} + \tilde{K}_t (y_t + v_t^{(k)} - g(\tilde{x}_t^{(k)}))$
  end for
end for

3 Using sequential Monte Carlo methods in the nimble package

This section describes how to specify a statistical model in the BUGS language and manipulate that model using the **nimble** package within the R (R Core Team

2015) statistical programming language. After describing some of the tools that **nimble** provides to interact with models in Section 3.1, we demonstrate the available SMC methods within nimble. Section 3.2 describes **nimble**'s SMC methods for inference in models with fixed parameters. Section 3.3 describes available SMC methods for state-space models with no fixed parameters. A supplement to the paper includes a full R script of all code shown below.

3.1 Creating and manipulating BUGS models within nimble

The **nimble** package uses the BUGS language to specify hierarchical statistical models. We will not describe model specification in the BUGS language here – interested readers can find a brief overview of writing BUGS models in the **nimble** User Manual (**nimble** Development Team 2016), or a more detailed guide in Lunn *et al.* (2012). Instead, we focus on how to interact with BUGS models using **nimble**. To introduce **nimble**'s features, we will use a linear Gaussian state-space model in which all parameters are fixed. Such a model will allow us to validate the results of our sequential Monte Carlo algorithms with the analytic solutions provided by the Kalman filter. Let y_t be the observed data at time point t, let x_t be the latent state at time point t, and suppose we have 10 time points. The model is:

$$\begin{aligned} x_1 &\sim N(0, 1) \\ x_t &\sim N(0.8 * x_{t-1}, 1) \quad for \quad t = 2, \ldots, 10 \\ y_t &\sim N(x_t, 0.5) \quad for \quad t = 1, \ldots, 10 \end{aligned}$$

where $N(\mu, \sigma^2)$ denotes the Normal distribution with mean μ and variance σ^2. We remark that although this example model is relatively simple, the algorithms presented in Sections 3.2 and 3.3 can be applied to any state-space model written in BUGS.

Models written in the BUGS language are read by the `nimbleCode` function. For example, BUGS code for the linear Gaussian model can be written and read into the **nimble** package as follows:

```
R> library('nimble')
R> set.seed(1)
R> exampleCode <- nimbleCode({
+    x[1] ~ dnorm(0, var = 1)
+    y[1] ~ dnorm(x[1], var = .5)
+    for(t in 2:10){
+      x[t] ~ dnorm(.8 * x[t-1], var = 1)
+      y[t] ~ dnorm(x[t], var = .5)
+    }
+  })
```

Once the model has been specified in BUGS, a `model` object can be created using the `nimbleModel` function. We provide `data` to the `nimbleModel` function in the form of a named list. In this example, we use simulated data from our model that is stored in a vector named `modelData`.

```
R> exampleModel <- nimbleModel(code = exampleCode, data = list(y = modelData))
```

3.2 Filtering given fixed parameters

Now that we have built a `model` object for the example model, we can use algorithms from **nimble**'s library to conduct inference. These algorithms are all written as functions within **nimble**'s DSL. We begin by demonstrating the use of the bootstrap filter (Section 2.3) to estimate the filtering distribution $p(x_t|y_{1:t})$.

```
R> exampleBootstrapFilter <- buildBootstrapFilter(exampleModel, nodes = 'x',
+    control = list(saveAll = TRUE, thresh = .9))
```

The `buildBootstrapFilter` function builds a bootstrap filter for the model given in the first argument. The `nodes` argument gives the name (or names) of the latent states to be filtered. Importantly, **nimble** filters require the latent states to have the same dimension at each time point. The algorithm parameters, packaged in the control list, include `saveAll` (should filtered state estimates be saved from all time points, or from just the last one) and `thresh` (a threshold for resampling, labeled τ in Section 2.4). Additional arguments to the control list can be found by calling `help(buildBootstrapFilter)`.

After the bootstrap filter has been built for the example model, it can be run in R by calling the `run` method of the filter, taking the number of particles to use as an argument, and returning an estimate of the log likelihood of the data.

```
R> exampleBootstrapFilter$run(100)

[1] -15.3636
```

For users wishing to write their own algorithms, constructing and running **nimble** functions in R allows for easy testing and debugging of algorithm logic. Once an algorithm has been successfully constructed in R, it can be compiled into C++ for efficient execution. Below, we compile the bootstrap filter algorithm using the `compileNimble` function, and run the compiled filter using 10,000 particles. Note that the model must be compiled before or in the same step as the algorithm.

```
R> CexampleModel <- compileNimble(exampleModel)
R> CexampleBootstrapFilter <- compileNimble(exampleBootstrapFilter,
+    project = exampleModel)
R> CexampleBootstrapFilter$run(10000)
R> bootstrapFilterSamples <- as.matrix(CexampleBootstrapFilter$mvEWSamples)
```

The bootstrap filter, like most filters in **nimble**, saves two arrays with samples from the filtering distribution. One array, named `mvEWSamples`, contains equally weighted samples from the filtering distribution. The second array, `mvWSamples`, contains non-equally weighted samples from the filtering distribution along with weights for each sample. These arrays can be easily converted to R matrices via the `as.matrix` function, which is used above to create the `bootstrapFilterSamples` object.

Next, we demonstrate **nimble**'s auxiliary particle filter algorithm (Section 2.4). The auxiliary particle filter is constructed in **nimble** similarly to the bootstrap filter, using a call to the `buildAuxiliaryFilter` function. **nimble**'s auxiliary filter allows users to choose between two lookahead functions: one that uses a simulation from the transition equation $\tilde{x}^{(k)}_{t|t-1} \sim f(x_t|x^{(k)}_{t-1})$, and one that uses the expected value of the transition equation $\tilde{x}^{(k)}_{t|t-1} = E(x_t|x^{(k)}_{t-1})$, via the `lookahead` control list argument.

```
R> exampleAuxiliaryFilter <- buildAuxiliaryFilter(exampleModel, nodes = 'x',
+    control = list(saveAll = TRUE, lookahead = 'mean'))
R> CexampleAuxiliaryFilter <- compileNimble(exampleAuxiliaryFilter,
+    project = exampleModel, resetFunctions = TRUE)
R> CexampleAuxiliaryFilter$run(10000)
R> auxiliaryFilterSamples <- as.matrix(CexampleAuxiliaryFilter$mvEWSamples)
```

The final method we demonstrate for models with fixed parameters is the ensemble Kalman filter, which can be built via a call to `buildEnsembleKF`. Note that the ensemble Kalman filter, as described in Section 2.7, does not produce weights with its particle estimates. Thus only one output array, named `mvSamples`, is available for this algorithm.

```
R> exampleEnsembleKF <- buildEnsembleKF(exampleModel, nodes = 'x',
+    control = list(saveAll = TRUE))
R> CexampleEnsembleKF <- compileNimble(exampleEnsembleKF,
+    project = exampleModel, resetFunctions = TRUE)
R> CexampleEnsembleKF$run(10000)
R> EnKFSamples <- as.matrix(CexampleEnsembleKF$mvSamples)
```

Since our example model has normal transition and observation equations, the filtering distribution can also be calculated analytically using the Kalman filter (Kalman 1960). Below, we use the **dlm** package (Petris 2010) to apply a Kalman filter to our model and compare the analytic filtering distribution provided by the Kalman filter to the approximate filtering distributions given by the bootstrap filter, auxiliary particle filter, and EnKF. Note that the quantiles in Figure 1 align almost exactly for all filters.

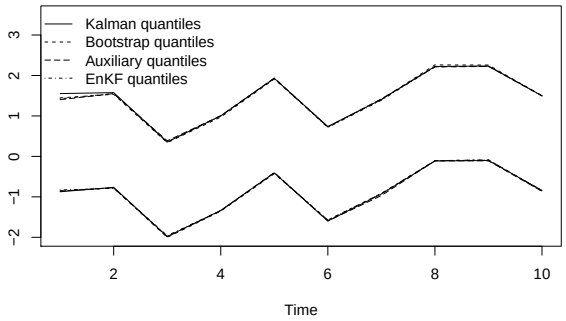

Figure 1: 2.5 % and 97.5 % quantiles of the filtering distribution for the Kalman filter and nimble's particle filters.

3.3 Inference on models with unknown parameters

The example model in the previous section had no unknown parameters – an uncommon scenario for dealing with real data. We next demonstrate **nimble**'s Liu and West filter and PMCMC algorithms, both of which can be used to estimate the posterior distributions of unknown top-level parameters in state-space models. To demonstrate these algorithms, we first construct a stochastic volatility model, which we use to model latent volatility in daily exchange rates.

We use the stochastic volatility model outlined in Pitt and Shephard (1999a). Let r_t be the exchange rate at time t, and define y_t as 100 times the daily log return, that is, $y_t = 100 \times (log(r_t) - log(r_{t-1}))$ for $t = 2, \ldots, T$. Our stochastic volatility model is then

$$
\begin{aligned}
y_t &= \varepsilon_t \beta \exp(\frac{x_t}{2}), \quad & \varepsilon_t &\sim N(0,1) \\
x_t &= \phi x_{t-1} + \nu_t, \quad & \nu_t &\sim N(0, \sigma^2)
\end{aligned}
$$

In this model, β can be interpreted as the constant volatility, while x_t is the latent, evolving volatility. Following Pitt and Shephard (1999b), prior distributions are placed on the parameters β, ϕ, and σ as follows:

$$
\begin{aligned}
\phi^* &\sim B(18,1), \quad & \phi &= 2\phi^* - 1 \\
\sigma^2 &\sim IG(5, \frac{1}{20}) \\
\beta^2 &\sim IG(5, \frac{1}{20})
\end{aligned}
$$

Above, $B(a, b)$ denotes the beta distribution with parameters a and b, and $IG(c, d)$ denotes the inverse gamma distribution with shape parameter c and scale parameter d. The stochastic volatility model can be written in BUGS code as

```
R> stochVCode <- nimbleCode({
+    x[1] ~ dnorm(phi * x0, sigmaSquaredInv)
+    y[1] ~ dnorm(0, var = betaSquared * exp(x[1]))
+    for(t in 2:T){
+      x[t] ~ dnorm(phi * x[t-1],  sigmaSquaredInv)
+      y[t] ~ dnorm(0, var = betaSquared * exp(x[t]))
+    }
+
+    x0 ~ dnorm(1, sigmaSquaredInv)
+    phi <- 2 * phiStar - 1
+    phiStar ~ dbeta(18, 1)
+    sigmaSquaredInv ~ dgamma(5, 20)
+    betaSquared <- 1 / betaSquaredInv
+    betaSquaredInv ~ dgamma(5, 20)
+  })
```

Note that above the precision parameterization of the normal distribution is used for the latent volatility x.

We use as data exchange rates for the U.S. Dollar (USD) quoted in Euros (EUR) starting on January 1st, 2012, and continuing for 66 days after that. This data set can be found in the **stochvol** R package (Kastner 2016). The **stochvol** package also includes a `logret` function to calculate log returns.

```
R> library('stochvol')
R> data('exrates')
R> y <-  100 * logret(exrates$USD[exrates$date > '2012-01-01'])
```

We next create and compile a `model` object for the above BUGS code, again following Pitt and Shephard (1999a) by using as starting values $\beta = .5992$, $\phi = .9702$, $\sigma = .178$, and providing T as a constant.

```
R> stochVolModel <- nimbleModel(code = stochVCode, name = 'stochVol',
+    constants = list(T = 67),  data = list(y = y),
+    inits = list(betaSquaredInv = 2.785,  phi = .9702,
+      sigmaSquaredInv = 31.561))
R> CstochVolModel <- compileNimble(stochVolModel)
```

To build a Liu and West filter (as detailed in Section 2.5), we use the `buildLiuWestFilter` function. The Liu and West filter requires specification not only of the latent states (via the `nodes` argument), but also of the top level parameters to be estimated (via the `params` argument). Additionally, the Liu and West filter does not return an estimate of the log likelihood of the data.

```
R> stochVolLiuWestFilter <- buildLiuWestFilter(model = stochVolModel,
+    nodes = 'x', params = c( 'betaSquaredInv', 'phiStar',
+      'sigmaSquaredInv'))
R> CstochVolLiuWestFilter <- compileNimble(stochVolLiuWestFilter,
+    project = stochVolModel)
R> CstochVolLiuWestFilter$run(50000)
```

Once the Liu and West filter has been run, we can extract the posterior distribution of top-level parameters. The code below creates a histogram of the posterior distribution of the σ parameter.

```
R> sigmaSquaredSamples <- 1 / as.matrix(CstochVolLiuWestFilter$mvEWSamples,
+    'sigmaSquaredInv')
R> hist(sigmaSquaredSamples, main ='', xlab = '')
```

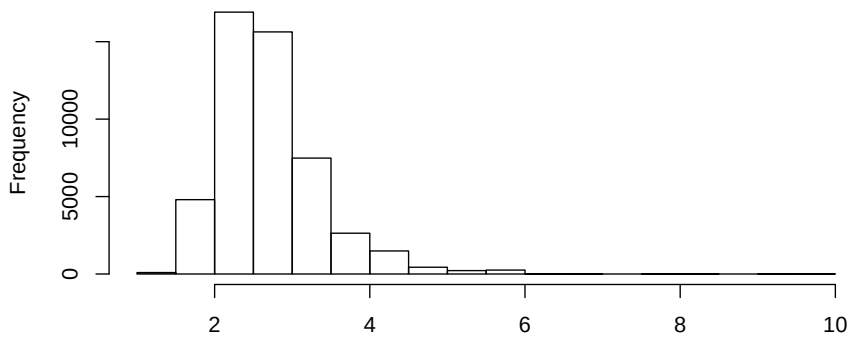

Figure 2: Histogram of the posterior distribution of σ^2 from the Liu and West filter.

An alternative set of methods for inference on top-level parameters in state-space models is Particle MCMC (Section 2.6). **nimble**'s Particle Marginal Metropolis Hastings (PMMH) sampler takes advantage of **nimble**'s existing MCMC framework, in which specific samplers can be specified for different nodes in a model. For a full description of **nimble**'s MCMC capabilities, reference Chapter 7 of the **nimble** User Manual (**nimble** Development Team 2016).

The PMMH sampler in **nimble** uses a normal proposal distribution, and can be applied to sample either scalar parameters (using the `RW_PFilter` sampler) or vectors of parameters (using the `RW_PFilter_block` sampler). To implement the PMMH algorithm, we first set up an MCMC specification for our stochastic volatility model using the `configureMCMC` function. The PMMH sampler can

be added to the MCMC specification with a call to the `addSampler` function. Additional options to customize the sampler can be specified within the `control` list.

```
R> stochVolMCMCSpec <- configureMCMC(stochVolModel, nodes = NULL,
+    monitors = c('betaSquaredInv', 'phi', 'sigmaSquaredInv', 'x'), thin = 20)
R> stochVolMCMCSpec$addSampler(target = c('betaSquaredInv', 'phiStar',
+    'sigmaSquaredInv'), type = 'RW_PF_block',
+    control = list(propCov = .1 * diag(3),
+    adaptive = TRUE, latents = 'x', pfResample = TRUE))
```

The `control` list argument is used to set the initial proposal covariance, here specified to be a 3×3 diagonal matrix with 0.1 entries on the diagonal. Additionally, the PMMH sampler can be set to use an adaptive algorithm to tune the proposal covariance matrix as the algorithm runs via the `adaptive` argument. The `resample` argument allows the algorithm to resample $\tilde{p}(y_{1:T}|\theta^{i-1})$ at the beginning of each iteration (before Step 4 in Algorithm 4). This can help to reduce the chance that the algorithm gets "stuck" at a particular set of parameters due to a high marginal likelihood estimate. High likelihood estimates can arise naturally because of the stochastic nature of particle filter likelihood estimation.

Once the PMMH sampler is added to the MCMC specification, the algorithm can be built using the `buildMCMC` function, and then compiled. Posterior samples are stored in `cMCMC$mvSamples`. Below we demonstrate running **nimble**'s PMMH algorithm for 10,000 iterations and extracting posterior samples of the β^{-2} parameter. Note that the first 100 thinned samples of our MCMC output are discarded as a burn-in period.

```
R> stochVolMCMC <- buildMCMC(stochVolMCMCSpec)
R> cMCMC <- compileNimble(stochVolMCMC, project = stochVolModel,
+    resetFunctions = TRUE)
R> cMCMC$run(10000)
R> mcmcOut <- as.matrix(cMCMC$mvSamples, 'betaSquaredInv')[-c(1:100),]
```

The **coda** package provides tools for analyzing MCMC output (Plummer *et al.* 2006). The code below creates a trace plot and posterior density plot for the β^{-2} parameter in our model.

```
R> library('coda')
R> mcmcOut <- as.mcmc(mcmcOut)
R> traceplot(mcmcOut)
R> densplot(mcmcOut)
```

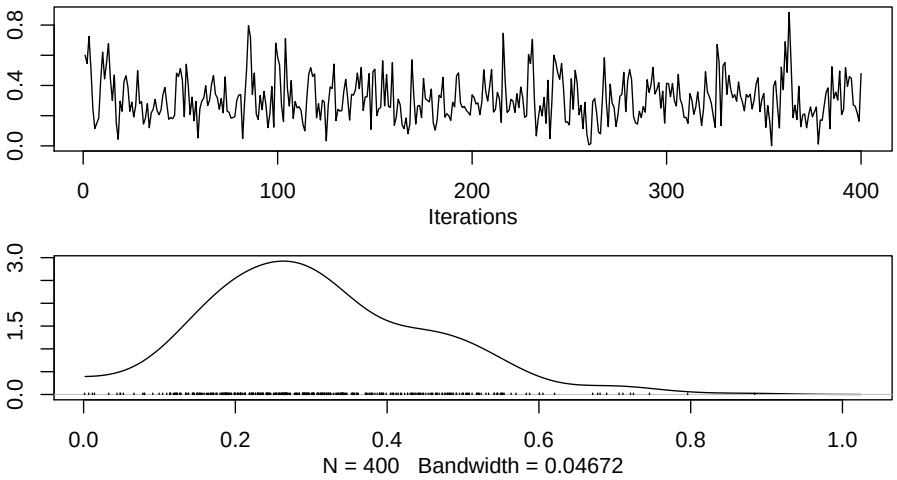

Figure 3: Trace plot and posterior density plot for β^{-2} from PMMH sampler.

4 Programming SMC algorithms in nimble

In this section, we demonstrate how **nimble** can be used to program model-generic SMC algorithms, that is, algorithms that can be applied to any state-space model written in BUGS. We will not cover every detail of nimbleFunction programming, but rather wish to show how algorithms are expressed compactly in high-level code that gets compiled via C++. This will allow other programmers to quickly adapt our functions for their own needs. For a more detailed discussion of nimbleFunction programming, see de Valpine *et al.* (in-press) or Chapter 9 of **nimble** Development Team (2016).

We demonstrate programming in **nimble** by providing code for a bootstrap filtering algorithm. Note that the code shown below is simpler than the actual implementation of the bootstrap filter available in **nimble** through the `buildBootstrapFilter` function. However, this demonstration code is indeed a fully functional bootstrap filter – the bootstrap filter included in the **nimble** package simply has more customization options than the demonstration algorithm provided here.

Each function written in **nimble** has two different types of code: `setup` code and `run` code. When a function is called in **nimble**, the `setup` code is evaluated first. `setup` code is written in R and is primarily used to extract model information for later use in the `run` code. After `setup` code has been used to prepare the algorithm, `run` code is executed. `run` code is written in the **nimble** DSL, which allows the code to be compiled into C++, in turn providing efficient execution of an algorithm's computations. `run` code can make use of objects created in the `setup` code.

The first function shown below, named `bootstrapFilter`, extracts necessary

information from the state-space model, including the names and dimensions of the latent nodes to be sampled. It then iterates through time, at each time point t calling the second function, `bootstrapStep`. The `bootstrapStep` function takes information about the latent state at time t, and then conducts Steps 7 through 13 of the bootstrap filter algorithm given in Section 2.3. As the filtering algorithm progresses through each time point, samples from the filtering distribution at that time are saved in **nimble** `modelValues` objects. `modelValues` objects provide containers for storing values of model nodes in **nimble**.

Below is the call to `nimbleFunction` that defines `setup` and `run` code for the `bootstrapFilter`. Note that the `setup` code for the `myBootstrapFilter` function is provided with the **nimble** model object and the names of the latent states as arguments. The setup code first defines a function that will initialize the model and then obtains the names and dimensions of the latent states in the model in time order. Two `modelValues` objects are created to store samples from the latent states. The `mvWSamples` object will store non-equally weighted samples, while `mvEWSamples` will store equally weighted samples. Finally, the `setup` code creates a list of `bootstrapStep` functions (called a `nimbleFunctionList`). For each time point $t = 1, \ldots, T$, the list contains one `bootstrapStep` function, which will conduct a bootstrap filtering algorithm at that time. We note that the creation of a separate `bootstrapStep` function for each time point t is necessary to allow the latent state x_t at each time point to have potentially different observation dependencies y_t.

```
R> bootstrapFilter <- nimbleFunction(
+    setup = function(model, latentNodes) {
+      my_initializeModel <- initializeModel(model)
+      latentNodes <- model$expandNodeNames(latentNodes, sort = TRUE)
+      dims <- lapply(latentNodes, function(n) nimDim(model[[n]]))
+      mvWSpec <- modelValuesConf(vars = c('x', 'wts'),
+        types = c('double', 'double'),
+        sizes = list(x = dims[[1]], wts = 1))
+      mvWSamples <- modelValues(mvWSpec)
+      mvEWSpec <- modelValuesConf(vars = c('x'), types = c('double'),
+        sizes = list(x = dims[[1]]))
+      mvEWSamples <- modelValues(mvEWSpec)
+      bootStepFunctions <- nimbleFunctionList(bootstrapStepVirtual)
+      timePoints <- length(latentNodes)
+      for (t in 1:timePoints)
+        bootStepFunctions[[t]] <- bootstrapStep(model, mvWSamples,
+          mvEWSamples, latentNodes, t)
+    },
+    run = function(K = integer()) {
+      my_initializeModel$run()
+      resize(mvWSamples, K)
+      resize(mvEWSamples, K)
+      for (t in 1:timePoints)
```

```
+         bootStepFunctions[[t]]$run(K)
+     }
+ )
```

The `run` code for the `bootstrapFilter` function takes as its only input argument the number of particles (K) to use for estimation. `run` code requires explicit specification of the type of any input arguments, so here K is specified as an integer object. In general, the type of object to be returned must also be specified, although this function does not return any objects so no specification is necessary. The `run` function first initializes the model (conducting Steps 2 and 3 of the bootstrap filter algorithm), and then re-sizes the `modelValues` objects so that they can store K particles. After that, the `run` function iterates through each time point, running the `bootstrapStep` function that was defined for that time point in the `setup` code. Note that this example algorithm does not provide an estimate of the likelihood $\tilde{p}(y_{1:T})$.

Creating a `nimbleFunctionList`, such as the one used in the `setup` code above, requires an additional piece of code that informs **nimble** about the input arguments and return objects of each function in that list. Specifically, the `nimbleFunctionVirtual` function is used to define the attributes that each function in the `nimbleFunctionList` will have. Below, we specify that each element of our `nimbleFunctionList` will have a run function with a single integer input.

```
R> bootstrapStepVirtual <- nimbleFunctionVirtual(
+     run = function(K = integer()) {}
+ )
```

`setup` and `run` code for the `bootstrapStep` function are given below. At each time point t, the `setup` function gets the names and deterministic dependencies of the previous and current latent states. The `run` code first declares a length K vector of integers (to store particle indices) and a length K vector of doubles (to store particle weights). The `run` code then iterates through the particles. For each particle, the code takes the value of the latent state at $t-1$ from the equally weighted `modelValues` object, uses that value to propagate a value for the latent state at time t, and calculates a weight. The particles and corresponding weights are stored in the non-equally weighted `modelValues` object. Finally, particles are resampled proportional to their weights and the resampled particles are stored in the equally weighted `modelValues` object.

In this algorithm particles are propagated using the proposal distribution $q(x_t|x_{t-1}^{(k)}, y_t) = f(x_t|x_{t-1}^{(k)})$ which simplifies the weight calculation in Step 8 of Algorithm 1. Additionally, since resampling is performed at each time point, weights from time $t-1$ do not need to be used when calculating weights at time t. This results in a weight calculation of $w_t^{(k)} = g(y_t|\tilde{x}_t^{(k)})$.

```
R> bootstrapStep <- nimbleFunction(
+     contains = bootstrapStepVirtual,
```

```
+    setup = function(model, mvWSamples, mvEWSamples, latentNodes,
+      timePoint) {
+      notFirst <- timePoint != 1
+      prevNode <- latentNodes[if(notFirst) timePoint - 1 else timePoint]
+      thisNode <- latentNodes[timePoint]
+      prevDeterm <- model$getDependencies(prevNode, determOnly = TRUE)
+      thisDeterm <- model$getDependencies(thisNode, determOnly = TRUE)
+      thisData   <- model$getDependencies(thisNode, dataOnly = TRUE)
+    },
+    run = function(K = integer()) {
+      ids <- integer(K, 0)
+      wts <- numeric(K, 0)
+      for(k in 1:K) {
+        if(notFirst) {
+          copy(from = mvEWSamples, to = model, nodes = 'x',
+            nodesTo = prevNode, row = k)
+          calculate(model, prevDeterm)
+        }
+        simulate(model, thisNode)
+        copy(from = model, to = mvWSamples, nodes = thisNode,
+          nodesTo = 'x', row = k)
+        calculate(model, thisDeterm)
+        wts[k] <- exp(calculate(model, thisData))
+        mvWSamples['wts', k][1] <<- wts[k]
+      }
+      rankSample(wts, K, ids)
+      for(k in 1:K){
+        copy(from = mvWSamples, to = mvEWSamples, nodes = 'x',
+          nodesTo = 'x', row = ids[k], rowTo = k)
+      }
+  })
```

The calls to `calculate` within the above `run` code serve two purposes. The first two `calculate` calls are used to calculate the values of any deterministic dependencies of the latent state, as these dependencies must be recalculated any time the latent state takes on a new value. The third call to `calculate` is used to calculate the log-likelihood of the data given the current latent state value, which is then used as a particle weight. Note that the `rankSample` function fills the elements of the `ids` vector with the indices of the particles that have been chosen in the resampling procedure.

Once the `nimbleFunction`s have been defined, we can build, compile, and run the bootstrap filter. The code below runs the example filter on the `exampleModel` of Section 3.2 and creates a histogram of samples from the filtering distribution of x at the last time point.

```
R> myBootstrap <- bootstrapFilter(exampleModel, 'x')
R> cmyBootstrap <- compileNimble(myBootstrap, project = exampleModel,
```

```
+    resetFunctions = TRUE)
R> cmyBootstrap$run(1000)
R> filterSamps <- as.matrix(cmyBootstrap$mvEWSamples, 'x')
R> hist(filterSamps, main = '', xlab = '')
```

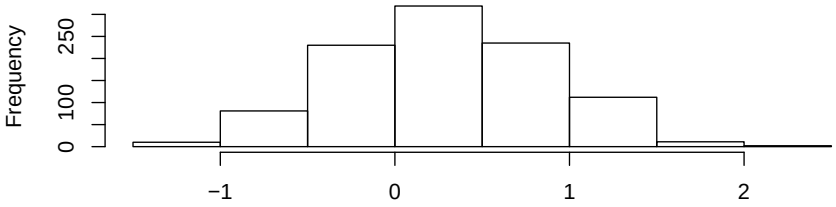

Figure 4: A histogram of the filtering distribution of x.

The bootstrap filter code provided above demonstrates **nimble**'s ability to program model-generic algorithms. The filter could be used to conduct filtering on any correctly specified state-space model. In addition to the generality of the algorithm, it would be relatively straightforward to modify the filter, changing it to an auxiliary particle filter, a Liu and West filter, or a filter type not currently included in **nimble**. The ease with which existing algorithms can be modified, along with the generality with which they are written, promotes the development of user-written filters in a manner previously unavailable in R or other software environments.

5 Conclusion

This paper has described **nimble**'s suite of SMC algorithms, which provide a straightforward method of conducting inference on state-space models. In addition, **nimble**'s model-generic programmability make it perfectly suited for implementing new SMC algorithms, an example of which was given in Section 4. **nimble**'s flexible model specification also enables the application of existing algorithms to models that do not fall into the traditional state-space model framework. For example, a model could be specified where a number of state-space models are set within a larger hierarchical structure. Using **nimble**, SMC algorithms could be used to estimate the individual state-space models, while an MCMC algorithm could conduct inference on higher-level parameters.

Additional examples of modeling and inference using **nimble** can be found at `http://r-nimble.org/`.

Acknowledgements

This work was supported by the NSF under grant DBI-1147230, and by the Google Summer of Code.

References

Andersson H, Britton T (2000). *Stochastic Epidemic Models and Their Statistical Analysis*, volume 151. Springer Science & Business Media.

Andrieu C, Doucet A, Holenstein R (2010). "Particle Markov Chain Monte Carlo Methods." *Journal of the Royal Statistical Society B*, **72**(3), 269–342.

Carpenter J, Clifford P, Fearnhead P (1999). "Improved Particle Filter for Nonlinear Problems." In *IEE Proceedings – Radar, Sonar and Navigation*, volume 146, pp. 2–7. IET.

Carvalho C, Johannes MS, Lopes HF, Polson N (2010). "Particle Learning and Smoothing." *Statistical Science*, **25**(1), 88–106.

Creal D (2012). "A Survey of Sequential Monte Carlo Methods for Economics and Finance." *Econometric Reviews*, **31**(3), 245–296. doi:10.1080/07474938.2011.607333. `http://dx.doi.org/10.1080/07474938.2011.607333`, URL `http://dx.doi.org/10.1080/07474938.2011.607333`.

de Valpine P, Turek D, Paciorek CJ, Anderson-Bergman C, Lang DT, Bodik R (in-press). "Programming with Models: Writing Statistical Algorithms for General Model Structures with **nimble**." *Journal of Computational and Graphical Statistics*. doi:10.1080/10618600.2016.1172487. `http://dx.doi.org/10.1080/10618600.2016.1172487`, URL `http://dx.doi.org/10.1080/10618600.2016.1172487`.

Doucet A, De Freitas N, Gordon N (2001). "An Introduction to Sequential Monte Carlo Methods." In *Sequential Monte Carlo Methods in Practice*, pp. 3–14. Springer.

Doucet A, Godsill S, Andrieu C (2000). "On Sequential Monte Carlo Sampling Methods for Bayesian Filtering." *Statistics and Computing*, **10**(3), 197–208.

Doucet A, Johansen AM (2009). "A Tutorial on Particle Filtering and Smoothing: Fifteen Years Later." In *Handbook of Nonlinear Filtering*, pp. 656–704. Oxford University Press.

Evensen G (2003). "The Ensemble Kalman Filter: Theoretical Formulation and Practical Implementation." *Ocean Dynamics*, **53**, 343–367.

Gillijns S, Mendoza OB, Chandrasekar J, De Moor B, Bernstein D, Ridley A (2006). "What Is the Ensemble Kalman Filter and How Well Does It Work?" In *American Control Conference, 2006*, pp. 6–pp. IEEE.

Gordon NJ, Salmond DJ, Smith AF (1993). "Novel Approach to Nonlinear/Non-Gaussian Bayesian State Estimation." In *IEE Proceedings F*, volume 140, pp. 107–113. IET.

Kalman RE (1960). "A New Approach to Linear Filtering and Prediction Problems." *Transactions of the ASME, Journal of Basic Engineering*, **82**(Series D), 35–45.

Kastner G (2016). "Dealing with Stochastic Volatility in Time Series Using the R Package **stochvol**." *Journal of Statistical Software*, **69**(1), 1–30. ISSN 1548-7660. doi:10.18637/jss.v069.i05. URL `https://www.jstatsoft.org/index.php/jss/article/view/v069i05`.

King A, Nguyen D, Ionides E (2016). "Statistical Inference for Partially Observed Markov Processes via the R Package **pomp**." *Journal of Statistical Software*, **69**(1), 1–43. ISSN 1548-7660. doi:10.18637/jss.v069.i12. URL `https://www.jstatsoft.org/index.php/jss/article/view/v069i12`.

Knape J, de Valpine P (2012). "Fitting Complex Population Models by Combining Particle Filters with Markov Chain Monte Carlo." *Ecology*, **93**(2), 256–263. ISSN 1939-9170. doi:10.1890/11-0797.1. URL `http://dx.doi.org/10.1890/11-0797.1`.

Liu J, West M (2001). "Combined Parameter and State Estimation in Simulation-Based Filtering." In *Sequential Monte Carlo Methods in Practice*, pp. 197–223. Springer.

Lunn D, Jackson C, Best N, Thomas A, Spiegelhalter D (2012). *The BUGS Book: A Practical Introduction to Bayesian Analysis*. Chapman & Hall/CRC Texts in Statistical Science. Taylor & Francis. ISBN 9781584888499. URL `https://books.google.com/books?id=Cthz3XMa_VQC`.

Murray L (2015). "Using **LibBi**." *Journal of Statistical Software*, **67**(1), 1–36. ISSN 1548-7660. doi:10.18637/jss.v067.i10. URL `https://www.jstatsoft.org/index.php/jss/article/view/v067i10`.

nimble Development Team (2016). ***nimble*** *User Manual, Version 0.5-1*. URL `http://r-nimble.org`.

Petris G (2010). "An R Package for Dynamic Linear Models." *Journal of Statistical Software*, **36**(1), 1–16. ISSN 1548-7660. doi:10.18637/jss.v036.i12. URL `http://www.jstatsoft.org/index.php/jss/article/view/v036i12`.

Pitt MK (2002). "Smooth Particle Filters for Likelihood Evaluation and Maximisation."

Pitt MK, Shephard N (1999a). "Filtering via Simulation: Auxiliary Particle Filters." *Journal of the American Statistical Association*, **94**(446), 590–599.

Pitt MK, Shephard N (1999b). "Time Varying Covariances: a Factor Stochastic Volatility Approach." *Bayesian Statistics*, **6**, 547–570.

Plummer M, Best N, Cowles K, Vines K (2006). "**CODA**: Convergence Diagnosis and Output Analysis for MCMC." *R News*, **6**(1), 7–11. URL `http://CRAN.R-project.org/doc/Rnews/`.

Polson NG, Stroud JR, Müller P (2008). "Practical Filtering with Sequential Parameter Learning." *Journal of the Royal Statistical Society B*, **70**(2), 413–428.

R Core Team (2015). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL `http://www.R-project.org/`.

Smith A, Doucet A, de Freitas N, Gordon N (2001). *Sequential Monte Carlo Methods in Practice*. Springer Science & Business Media.

Storvik G (2002). "Particle Filters for State-Space Models with the Presence of Unknown Static Parameters." *IEEE Transactions on Signal Processing*, **50**(2), 281–289.

Todeschini A, Caron F, Fuentes M, Legrand P, Del Moral P (2014). "**Biips**: Software for Bayesian Inference with Interacting Particle Systems." *arXiv preprint arXiv:1412.3779*.

Wikle CK, Milliff RF, Herbei R, Leeds WB (2013). "Modern Statistical Methods in Oceanography: A Hierarchical Perspective." *Statistical Science*, **28**(4), 466–486. doi:10.1214/13-STS436. URL `http://dx.doi.org/10.1214/13-STS436`.

Zhou Y (2015). "**vSMC**: Parallel Sequential Monte Carlo in C++." *Journal of Statistical Software*, **62**(1), 1–49. ISSN 1548-7660. doi:10.18637/jss.v062.i09. URL `https://www.jstatsoft.org/index.php/jss/article/view/v062i09`.